Position Statement for
SC '21 panel: "The Computing Cambrian Explosion:
The Future of HPC with Non-Von Neumann Computing"

Sandia National Laboratories

Reversible Computing—The Long-Term Future of General Digital Computing

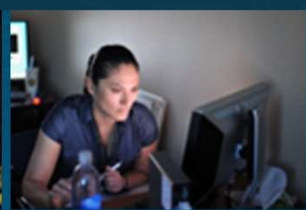

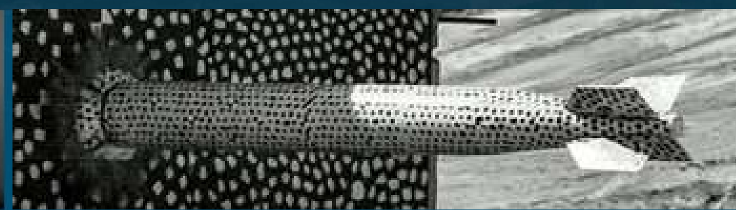

Thursday, November 18th, 2021

Michael P. Frank, Center for Computing Research

U.S. Department of Energy NNSA

Approved for public release, SAND2021-13358 C

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Reversible Computing (RC)— The Long-Term Future of General Digital Computing

From my POV: The adjective "***non-conventional***" would really be a more appropriate/inclusive descriptor than "non-von Neumann" for discussing forward-thinking future computing technologies.

- The qualifier "von Neumann" normally only refers to a class of (programmable) *processing architectures.*
  - Specifically, those with an instruction/data memory separated from an instruction-execution core (*e.g.*, typical CPUs, DSPs, GPUs…)
- But, we already have a wide range of *non*-von Neumann processing architectures in *conventional digital CMOS!*
  - *E.g.*, systolic arrays for signal-processing/linear algebra functions, FPGAs, Coarse-Grained Reconfigurable Arrays (CGRAs), digital implementations of spiking neural processing architectures (*e.g.*, IBM TrueNorth, Intel Loihi), *etc.*
- ***Reversible computing*** is the (only) non-conventional computing paradigm that offers us a new long-term scaling path by which we can improve the efficiency of ***all digital computing*** by *unboundedly-large factors.*
  - Please note that the scope of this includes ***both*** von Neumann and (digital) non-von Neumann processing architectures, as well as all possible kinds of special-purpose (fixed-function) digital circuits & systems…
    - *E.g.*, an FFT core for use in a DSP; or any other given digital IP block; or any digital ASIC
  - **Not often appreciated:** The scope of reversible computing also *includes* **communication** and **memory** functions (not just logic!).
    - But, whether processing and memory functions are separated (as in von Neumann architectures) is orthogonal to whether RC is used.

Key selling point for reversible computing:

- We can ***prove*** (easily) from standard physics and information theory, that RC is the *only* possible means by which energy efficiency *and cost efficiency* of general digital computing can continue improving *indefinitely* (*i.e.*, with no known fundamental limits).

Crux of von Neumann arch.

Processing Core

Inst./Data Memory

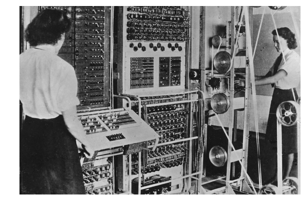

Colossus (1943) was non-von Neumann!

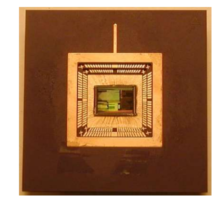

Pendulum ('99)

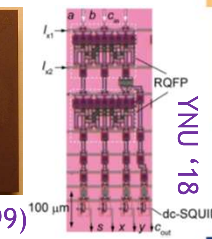

YNU '18

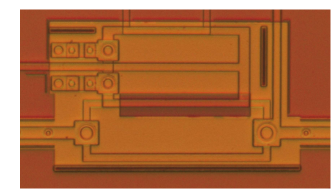

SFQ Rev. Mem. Cell ('20)

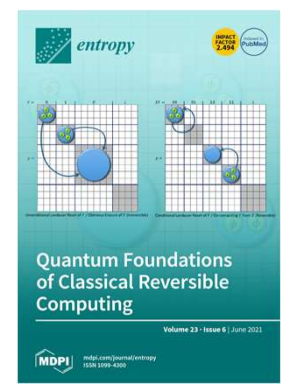

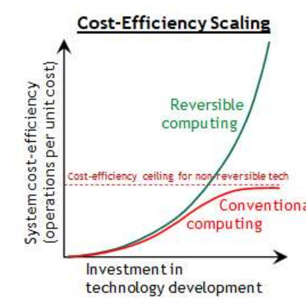

Answers to Panel Questions:

Key research challenge (for HPC applications):

- Package-level integration of high-quality integrated trapezoidal resonators (Sandia patent is pending).

Breakthrough performance for what applications?

- *All* computing applications, eventually (given ongoing manufacturing cost reductions).

Special-purpose or general-purpose?

- General.

How to use with existing programming approaches?

- Programming models & languages will need to evolve, first at low level (*e.g.*, ISA), then high-level.

How can software ease integration?

- EDA tools will need to be extended to support design of reversible hardware.

Manufacturing or production challenges?

- Yes; per-device manufacturing cost must continue declining in order for the overall cost-efficiency of reversible solutions to continue improving.
  - But, with RC we can stack logic in 3D w/o thermal issues!

Benefit vs. "drop-in" replacement of devices *etc.*

- In contrast to *all* physically possible refinements of conventional (*non*-reversible) computing, the long-term improvement possible w. RC is *not* fundamentally limited.

Tradeoffs in performance metrics:

- RC allows *any desired* level of improvement in ***energy*** efficiency (performance/watt) and *operating* cost-efficiency to be "bought"
  - in a tradeoff, through sufficient improvements in *manufacturing* cost-efficiency (and/or expected deployed system lifetimes).

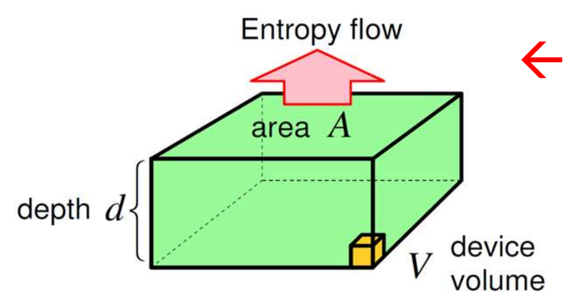

← RC max. perf./area ∝ $\sqrt{d}$

(Frank '97, "Ultimate Theoretical Models of Nanocomputers", doi:10.1088/0957-4484/9/3/005)